

VALID POST-CORRECTION INFERENCE FOR CENSORED REGRESSION PROBLEMS

BY YUEKAI SUN AND JONATHAN E. TAYLOR

Stanford University

Two-step estimators often called upon to fit censored regression models in many areas of science and engineering. Since censoring incurs a bias in the naive least-squares fit, a two-step estimator first estimates the bias and then fits a corrected linear model. We develop a framework for performing valid *post-correction inference* with two-step estimators. By exploiting recent results on post-selection inference, we obtain valid confidence intervals and significance tests for the fitted coefficients.

1. Introduction. Censored regression models were developed to handle otherwise ordinary statistical problems in which the data is subject to censoring. Formally, censored models are “two-part” models that adds to a linear model, e.g.

$$y = X\beta + \epsilon, \epsilon \sim \mathcal{N}(0, \sigma^2 I),$$

a “censoring” or “sample selection” step prior to observation $y = g(y^*)$. Perhaps, the most common example of a censored regression model is the standard (Type I) Tobit model ([Amemiya, 1985](#)):

$$(1.1) \quad y = \max\{0, y^*\}, y^* \sim \mathcal{N}(X\beta, \sigma^2 I).$$

The step $\max\{0, y^*\}$ may be changed to $\max\{y_0, y^*\}$ without essentially changing the model by absorbing y_0 into the intercept term. Applications of this model are found throughout econometrics and the social sciences. Two examples are family income after receiving welfare ([Keeley et al., 1978](#)) and the criminal behavior of habitual offenders ([Witte, 1980](#)).

Most previous work on censored regression models focuses on estimation. Two competing approaches are the maximum likelihood estimator (MLE) and various two-step estimators. Under appropriate regularity conditions, both approaches produce asymptotically normal estimates. To perform inference, we call upon the standard battery of likelihood ratio, score/Lagrange

MSC 2010 subject classifications: Primary 62N01; secondary 62P20

Keywords and phrases: censoring, sample selection, Tobit model, two-step estimator, post-correction inference, post-selection inference

multiplier, and Wald tests/intervals. Non-asymptotic inference is hard because the censoring mechanism incurs a bias in the usual least-squares estimate for β that must be corrected for. Two-step estimators account for the censoring bias by first estimating a correction term and then fitting a (linear) regression model with the correction term. Since the correction term is random, the usual confidence intervals/significance tests for the regression coefficients fail to have the nominal coverage/error rates.

We focus on performing inference with two-step estimators conditioned on the outcome of the censoring mechanism. Since the correction term is completely determined by the outcome of the censoring mechanism, we call this form of inference *post-correction inference*. Although we must condition on the censoring event to perform inference, the results are valid *unconditionally*. We discuss the unconditional validity of Section 4.

1.1. Background. The first censored regression model was proposed by Tobin (1958) to model household expenditure on durable goods. The model explicitly accounts for the fact that expenditure cannot be negative. Tobin called his model the model of *limited dependent variables*. It and its generalizations are commonly known in econometrics as *Tobit models*, a play on *probit models* by Goldberger (1964). Censored regression models are also broadly applicable in many other areas of science and engineering. For example, the survival times of patients and the time to failure of a machine or a system are both censored by the length of the study. In their respective domains, these models are called *survival* or *duration* models.

Since the 1970s, many generalizations of Tobin’s original model has appeared and continue to appear in econometrics and the aforementioned areas. Amemiya (1985) classified these models into five basic types according to the form of the likelihood function. We shall refer to various censored regression models according to Amemiya’s classification.

Fitting a censored regression model is usually performed by M-estimation. To fit simple models, the maximum likelihood estimator (MLE) (Amemiya, 1973) (usually after a change of variable proposed by Olsen (1978)) or Powell’s least absolute deviations (LAD) estimator (Powell, 1984) are preferred. The EM algorithm by Dempster, Laird and Rubin (1977) is also suitable for fitting censored regression models. For more complex models, two-step estimators (Heckman, 1976) are more computationally efficient. Heckman originally proposed his two-step estimator for a Type 3 Tobit model, but readily adapts to other censored regression models. Puhani (2000) compared two-step estimators with the MLE and concluded that two-step estimators are more robust when the covariates in the model are highly correlated.

Otherwise, the MLE is more (statistically) efficient.

Subject to appropriate regularity conditions, all the aforementioned estimators are asymptotically normal, thus inference is usually performed asymptotically. [Amemiya \(1973\)](#) showed the maximum likelihood estimate (MLE) is consistent and asymptotically normal (subject to homoscedasticity and normality). The same is generally true of two-step estimators. [Powell \(1984\)](#) showed his LAD estimator is consistent and asymptotically normal under more general conditions. This makes Powell’s estimator especially attractive when the data is nonnormal or the errors are heteroscedastic. The main drawback to Powell’s estimator is the computational expense of optimizing a nonsmooth function.

1.2. Related work on post-selection inference. This work is based on a framework for post-(model)-selection inference by [Lee et al. \(2013\)](#). At the core of both frameworks are key distributional results on the supremum of a Gaussian process restricted to a convex set $\mathcal{K} \subset \mathbf{R}^p$:

$$\sup_{\eta \in \mathcal{K}} \eta^T \epsilon, \quad \epsilon \sim \mathcal{N}(\mu, \Sigma).$$

In this paper, we give a geometric derivation of a special case (when the set \mathcal{K} is polyhedral) of the main result (Theorem 1) in [?](#) and refer to the source for a more general derivation. Like in [?](#) and [Lee et al. \(2013\)](#), normality is crucial to our results.

1.3. Notation. We follow the notation of [Amemiya \(1985\)](#). We are given $X \in \mathbf{R}^{n \times p}$ and responses $y \in \mathbf{R}^n$. We say x_i^T for the i -th row of X and y_i for the i -th response. Often we need to distinguish between the vectors and matrices of censored and all observations; the former appear with an underbar, e.g. \bar{y} is the vector of uncensored responses. Usually, the uncensored responses are constrained to fall in the non-negative orthant \mathbf{R}_+^n . Finally, $\phi(x)$ and $\Phi(x)$ are the pdf and CDF of the standard normal distribution.

2. Post-correction inference with two-step estimators.

2.1. Two two-step estimators. We describe two-step estimators in the context of the standard Tobit model. Since

$$\mathbf{E}[y_i \mid y_i > 0] = x_i^T \beta + \mathbf{E}[\epsilon_i \mid \epsilon_i > x_i^T \beta]$$

and the conditional expectation on the right side is usually nonzero (even when ϵ_i is not normal), performing (linear) regression on only the samples

with positive responses produces biased estimates. When ϵ is normal, the bias term simplifies to

$$(2.1) \quad \mathbf{E}[y_i \mid y_i > 0] = x_i^T \beta + \sigma \lambda \left(\frac{1}{\sigma} x_i^T \beta \right),$$

where $\lambda(x)$ is the inverse Mills ratio $\frac{\phi(x)}{1-\Phi(x)}$. A two-step estimator corrects the bias by first estimating the inverse Mills ratio $\lambda \left(\frac{1}{\sigma} x^T \beta \right)$ with probit regression, and then fitting a corrected linear model to the samples with positive responses. The corrected linear model is heteroscedastic:

$$(2.2) \quad \text{var}[\epsilon_i \mid \epsilon_i > x_i^T \beta] = \sigma^2 - \sigma^2 x_i^T \alpha \lambda \left(\frac{1}{\sigma} x_i^T \beta \right) - \sigma^2 \lambda \left(\frac{1}{\sigma} x_i^T \beta \right)^2,$$

so we can estimate β more efficiently with weighted least squares in the second step. Since our focus is not estimation, we skip this topic and refer to [Amemiya \(1985\)](#) for details.

Algorithm 1 Two-step estimator (Type 1 Tobit model)

Require: design matrix X , responses y

1: Estimate $\alpha = \frac{1}{\sigma} \beta$ with the probit MLE:

$$\hat{\alpha} = \arg \max_{\alpha} \prod_{i: y_i = 0} (1 - \Phi(x_i^T \alpha)) \prod_{i: y_i > 0} \Phi(x_i^T \alpha).$$

2: Regress the uncensored responses \bar{y} on \bar{X} and $\hat{\lambda} = \lambda(\bar{X} \hat{\alpha})$ to estimate β and σ :

$$\hat{\beta}, \hat{\sigma} = \arg \min_{\beta, \sigma} \frac{1}{2} \|\bar{y} - \sigma \hat{\lambda} - \bar{X} \beta\|_2^2.$$

[Heckman \(1979\)](#), who proposed the first two-step estimator, interprets the first step in the two-step estimator can be interpreted as obtaining a correction for the missing bias term in (2.1). As long as the probit MLE is consistent, the two-step estimator is also consistent. Further, when

1. $\{x_i^T\}$ are uniformly bounded,
2. $\lim_{n \rightarrow \infty} X^T X / n$ (exists and) is positive definite,
3. (β, σ) are in some (*a priori known*) compact subset of \mathbf{R}^p ,

$\sqrt{n}(\hat{\beta} - \beta)$ is asymptotically normal ([Amemiya, 1985](#)) and the usual battery of likelihood ratio, score, and Wald tests/intervals are valid. The two-step approach also generalizes readily to all five types of Tobit models in [Amemiya \(1985\)](#). We refer to [Amemiya \(1985\)](#) for details.

In the second step, one can also regress all observations (including the censored observations) to obtain an estimate of β . This estimator is also consistent and asymptotically normal. Which estimator is more efficient? Unfortunately, the answer depends on the parameter β . However, the framework also applies to this estimator by considering a *degenerate* constrained normal.

2.2. *A pivotal quantity for constrained normal variables.* Recall our goal is to perform inference on the fitted coefficients conditioned on the censoring event $\bar{y} > 0$. More generally, we are interested in targets of the form $\eta^T \bar{y}$ for some η . We begin by characterizing the conditional distribution of the response. Firstly, given $\bar{y} > 0$, \bar{y} has a constrained normal distribution; i.e. the conditional pdf of \bar{y} is $f(\bar{y} \mid \bar{y} > 0) \propto \phi(\bar{y}) \mathbf{1}_{\mathbf{R}_+^{\bar{n}}}(\bar{y})$. Secondly, given the censoring event, the correction term $\hat{\lambda}$ is simply

$$\hat{\lambda} = \lambda(\bar{X}\hat{\alpha}), \hat{\alpha} = \arg \max_{\alpha} \prod_{i \in S^c} (1 - \Phi(x_i^T \alpha)) \prod_{i \in S} \Phi(x_i^T \alpha).$$

The two-step estimate of $\gamma = (\beta, \sigma)$ is given by

$$\hat{\gamma} = \hat{Z}^\dagger \bar{y} = (\hat{Z}^T \hat{Z})^{-1} \hat{Z}^T \bar{y}, \hat{Z} = [\bar{X} \quad \hat{\lambda}].$$

To form (conditional) confidence intervals for or (conditionally) test the significance of the coefficients $\hat{\gamma}_j$, $j = 1, \dots, p$, we study its distribution conditioned on the censoring event. In this section, we exploit distributional results in [Lee et al. \(2013\)](#) to derive a pivotal quantity.

THEOREM 2.1. *Let $y \sim \mathcal{N}(\mu, \Sigma)$. Define $a = \frac{A\Sigma\eta}{\eta^T \Sigma \eta}$ and*

$$(2.3) \quad \mathcal{V}_\eta^+(y) = \sup_{j: a_j < 0} \frac{1}{a_j} (b_j - (Ay)_j + a_j \eta^T y)$$

$$(2.4) \quad \mathcal{V}_\eta^-(y) = \inf_{j: a_j > 0} \frac{1}{a_j} (b_j - (Ay)_j + a_j \eta^T y)$$

$$(2.5) \quad \mathcal{V}_\eta^0(y) = \inf_{j: a_j = 0} b_j - (Ay)_j.$$

Then, $\eta^T y$ conditioned on $Ay \leq b$ and $(\mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y)) = (v^+, v^-, v^0)$, has a truncated normal distribution, i.e.

$$\eta^T y \mid Ay \leq b, \mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y) \sim TN(\eta^T \mu, \eta^T \Sigma \eta, v^-, v^+).$$

We defer a proof to [Appendix A.1](#) and focus on a geometric interpretation of [Theorem 2.1](#). Let $C = \{y \in \mathbf{R}^n \mid Ay \leq b\}$. The functions $\mathcal{V}_\eta^+(y)$ and $\mathcal{V}_\eta^-(y)$ return the furthest we can move along $\tilde{\eta} = \Sigma^{1/2} \eta$ starting at $\tilde{y} = \Sigma^{-1/2} y$ while remaining in the set $\Sigma^{-1/2} C$, i.e.

$$\begin{aligned} \mathcal{V}_\eta^+(y) &= \sup\{\lambda \in \mathbf{R} : \lambda \tilde{\eta} + \tilde{y} \in \Sigma^{-1/2} C\} \\ \mathcal{V}_\eta^-(y) &= \inf\{\lambda \in \mathbf{R} : \lambda \tilde{\eta} + \tilde{y} \in \Sigma^{-1/2} C\}. \end{aligned}$$




FIG 1. Geometric intuition behind Theorem 2.1. To simplify the drawing, we assume $\Sigma = I$. By conditioning on $Ay \leq b, \mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y) \geq 0$, we are restricting ourselves to the dotted slice of the shaded region. Thus conditioning on $\mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y)$ is equivalent to conditioning on the component of y orthogonal to η . Since the orthogonal component is independent of $\eta^T y$, conditioning on it does not affect the normality of $\eta^T y$.

Since the set $\{\lambda : \lambda \tilde{\eta} + \tilde{y} \in \Sigma^{-1/2}C\}$ is equivalent to $\{\lambda : \lambda \Sigma^{1/2} \tilde{\eta} + \Sigma^{1/2} \tilde{y} \in C\}$, we have, in terms of the original variables, we have

$$\begin{aligned}\mathcal{V}_\eta^+(y) &= \sup\{\lambda : \lambda \Sigma \eta + y \in C\} \\ \mathcal{V}_\eta^-(y) &= \inf\{\lambda : \lambda \Sigma \eta + y \in C\}.\end{aligned}$$

The solution to both optimization problems are given by (2.3) and (2.4). By conditioning on $(\mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y)) = (v^+, v^-, v_0)$, we are further restricting y to fall on a slice of C (the slice $\{\lambda \Sigma \eta + y\} \cap C$ parametrized by λ). Theorem 2.1 merely says y restricted to this slice has a truncated normal distribution. This follows from the fact that $\mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y)$ are independent of $\eta^T y$ by construction.

Given $y \sim \mathcal{N}(\mu, \Sigma)$, we restrict to the slice $\{\lambda \Sigma \eta + y\} \cap \mathbf{R}_+^n$. To obtain a $\text{unif}(0, 1)$ distributed pivotal quantity, we apply a CDF transform. Figure 2 shows results from a simulation that empirically confirm the pivotal quantity is uniformly distributed.

COROLLARY 2.2. Let $y \sim \mathcal{N}(\mu, \Sigma)$. Define

$$(2.6) \quad F(x, \mu, \sigma^2, a, b) = \frac{\Phi\left(\frac{x-\mu}{\sigma}\right) - \Phi\left(\frac{v^--\mu}{\sigma}\right)}{\Phi\left(\frac{v^+-\mu}{\sigma}\right) - \Phi\left(\frac{v^--\mu}{\sigma}\right)}.$$




FIG 2. Empirical CDF of 10000 samples of the pivotal quantity (2.7) obtained by rejection sampling. The empirical distribution is very close to that of a $\text{unif}(0, 1)$ variable.

Then $F(\eta^T y, \eta^T \mu, \eta^T \Sigma \eta, \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^+(y))$ is a (conditional) pivotal quantity with a $\text{unif}(0, 1)$ distribution, i.e.

$$(2.7) \quad F(\eta^T y, \eta^T \mu, \eta^T \Sigma \eta, \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^+(y)) \mid Ay \leq b \sim \text{unif}(0, 1)$$

where $\mathcal{V}_\eta^+(y)$, $\mathcal{V}_\eta^-(y)$, $\mathcal{V}_\eta^0(y)$ are given by (2.3), (2.4), (2.5).

PROOF. By Theorem 2.1, we know

$$\eta^T y \mid y > 0, \mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y),$$

where $\mathcal{V}_\eta^+(y)$, $\mathcal{V}_\eta^-(y)$, $\mathcal{V}_\eta^0(y)$ are given by (2.3), (2.4), (2.5), has a truncated normal distribution. We apply the CDF transform to deduce

$$F(\eta^T y, \eta^T \mu, \eta^T \Sigma \eta, \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^+(y)) \mid y > 0, \mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y)$$

is uniformly distributed. Since this holds for any $\mathcal{V}_\eta^+(y)$, $\mathcal{V}_\eta^-(y)$, $\mathcal{V}_\eta^0(y)$, we deduce

$$F(\eta^T y, \eta^T \mu, \eta^T \Sigma \eta, \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^+(y)) \mid y > 0.$$

is also uniformly distributed. \square

2.3. Confidence intervals for the fitted coefficients. Recall we seek valid confidence intervals for β_j conditioned on the censoring event $\bar{y} > 0$. Let $\eta_j = (\bar{X}^\dagger)^T e_j$. By Corollary 2.2, we have

$$F(\eta_j^T \bar{y}, \eta_j^T \bar{X} \beta, \sigma^2 \|\eta_j\|_2^2, \mathcal{V}_{\eta_j}^-(\bar{y}), \mathcal{V}_{\eta_j}^+(\bar{y})) \mid \bar{y} > 0 \sim \text{unif}(0, 1),$$

where $\mathcal{V}_\eta^+(y)$, $\mathcal{V}_\eta^-(y)$, $\mathcal{V}_\eta^0(y)$ are given by (2.3), (2.4), (2.5) with $A = -I$ and $b = 0$. The second argument to F simplifies to

$$\eta_j^T \bar{X} \beta = e_j^T \bar{X}^\dagger \bar{X} \beta = \beta_j.$$

To obtain valid confidence intervals for β_j , we simply “invert” the pivotal quantity. The set

$$(2.8) \quad \left\{ \nu \in \mathbf{R} : \frac{\alpha}{2} \leq F(\eta_j^T \bar{y}, \nu, \sigma^2 \|\eta_j\|_2^2, \mathcal{V}_{\eta_j}^-(\bar{y}), \mathcal{V}_{\eta_j}^+(\bar{y})) \leq 1 - \frac{\alpha}{2} \right\}$$

is a $1 - \alpha$ confidence interval for β_j . The endpoints $\frac{\alpha}{2}$ and $1 - \frac{\alpha}{2}$ were chosen arbitrarily. By Lemma A.3, F decreases monotonically in ν . To obtain an intervals, we need to solve two univariate root-finding problem. Figure 3 shows results from two simulations that compare the coverage of the corrected intervals versus the normal intervals.

LEMMA 2.3. *Let $\eta_j = (\bar{X}^\dagger)^T e_j$. Define $\nu_\alpha(\bar{y})$ to be the (unique) root of*

$$F(\eta_j^T \bar{y}, \nu, \sigma^2 \|\eta_j\|_2^2, \mathcal{V}_{\eta_j}^-(\bar{y}), \mathcal{V}_{\eta_j}^+(\bar{y})) = \alpha.$$

Given $\bar{y} > 0$, $[\nu_{\alpha/2}(\bar{y}), \nu_{1-\alpha/2}(\bar{y})]$ is a valid $1 - \alpha$ confidence interval for β_j :

$$\mathbf{Pr}(\beta_j \in [\nu_{\alpha/2}(\bar{y}), \nu_{1-\alpha/2}(\bar{y})] \mid \bar{y} > 0) = 1 - \alpha.$$

The strength of the aforementioned results do come at a price. In this case, the price is wider confidence intervals. For example, suppose $\eta^T \bar{y}$ falls near an endpoint of the truncation interval $[v^-, v^+]$, say v^+ . Since this is likely to occur when $\eta^T \bar{\mu}$ is very large, there is a wide range of large values of $\eta^T \bar{\mu}$ that make this observation likely. However, when $\eta^T \bar{y}$ is not near an endpoint of the truncation interval (and the truncation interval) is large, then the confidence interval will be comparable to the least-squares (normal) intervals. Figure 4 compares the coverage of the truncated intervals with that of normal intervals.

2.4. *Testing the significance of fitted coefficients.* We also seek to test the significance of the j -th regression coefficient $\hat{\beta}_j$, i.e. test the hypothesis

$$H_0 : \eta_j^T \bar{X} \beta = \beta_j = 0, \eta_j = (\bar{X}^\dagger)^T e_j.$$

The output of the censoring mechanism is random, so H_0 is a *random* hypothesis. One possible interpretation of testing H_0 is we are testing H_0 conditioned on the censoring event. Since H_0 is fixed given $\bar{y} > 0$, this is




FIG 3. 95% corrected (left) and normal (right) intervals for β_1 in a Type 1 Tobit model over 50 simulated data sets ($n = 100$, $p = 10$, $\sigma^2 = 1.0$, standard Gaussian design). A green bar means the confidence interval covers β_1 while a red bar means otherwise. Although the corrected intervals are wider, they cover β_1 (black line) at the nominal rate. Empirically, we observe the coverage of the normal intervals becomes worse as the signal strength decreases.




FIG 4. 95% truncated (gray region) and normal (area between blue lines) confidence intervals as a function of $\frac{x}{\sigma}$. The truncation regions are $[-3\sigma, \infty]$ and $[-3\sigma, 3\sigma]$ in the figures on the left and right. When x falls near the boundary of the truncation region (in terms of σ), the conditional interval is much wider than the normal interval.

a hypothesis in the conventional sense. As we shall see, there is also an *unconditional* interpretation of testing H_0 .

Under H_0 , we know (by Corollary 2.2)

$$F(\eta_j^T \bar{y}, 0, \sigma^2 \|\eta_j\|_2^2, \mathcal{V}_{\eta_j}^-(\bar{y}), \mathcal{V}_{\eta_j}^+(\bar{y})) \mid \bar{y} > 0 \sim \text{unif}(0, 1),$$

We reject H_0 when the pivotal quantity falls in the top or bottom $\frac{\alpha}{2}$ quantile to obtain a valid α -level test for H_0 . This test is the post-correction counterpart to the usual t-test and controls Type I error conditioned on the censoring event:

$$\mathbf{Pr}(\text{reject } H_0 \mid \bar{y} > 0, H_0) \leq \alpha.$$

Since the test controls the Type I error rate at 5% for all possible outcomes of the censoring mechanism, the test also controls Type I error *unconditionally*. In other words, the test also is a valid unconditional test of H_0 .

LEMMA 2.4. *Let $\eta_j = (\bar{X}^\dagger)^T e_j$. Given $\bar{y} > 0$, the test that rejects when*

$$F(\eta_j^T \bar{y}, 0, \sigma^2 \|\eta_j\|_2^2, \mathcal{V}_{\eta_j}^-(\bar{y}), \mathcal{V}_{\eta_j}^+(\bar{y})) \in [0, \frac{\alpha}{2}) \cup (\frac{\alpha}{2}, 1]$$

is a valid α -level test for $H_0 : \beta_j = 0$.

3. Handling other censored regression models. Two-step estimators are broadly applicable to many censored regression models, and the post-correction inference framework also generalizes accordingly. We already described how to perform post-correction inference for the standard (Type 1) Tobit model in Section 2. In this section, we outline how the framework applies to other censored regression models.

3.1. Accelerated failure time models. Accelerated failure time (AFT) models are often used to model survival or duration data in science and engineering. Mathematically, AFT models are very similar to Tobit models. Let $t_i, i = 1, \dots, n$ be the (possibly censored) failure times of n units. AFT models posit a linear relationship between $\log(t_i)$ and some variables x_i^T ; i.e.

$$(3.1) \quad \log(t_i) = x_i^T \beta + \epsilon_i, \quad i = 1, \dots, n$$

where ϵ is independent error. The distribution of ϵ determines AFT model in use. Since we only observe t_i if unit i fails before the end of the testing period (at some known T), the failure times are right-censored at T . If the error is normally distributed (a log-normal AFT model), then we recognize the standard Tobit model with right-censoring at $\log(T)$:

$$(3.2) \quad y \sim \min\{y^*, \log(T)\}, \quad y^* \sim \mathcal{N}(X\beta, \sigma^2 I).$$

To handle right censoring at $\log(T)$, we absorb $\log(T)$ into the intercept term and take $-\log(t_i)$ as the response. If the testing period is different for each unit, the model is slightly changed because the resulting model is equivalent to (3.2) where an entry of β is known. To fit a log-normal AFT model, we simply take $\log(T_i) - \log(t_i)$ as the response and call upon Algorithm 1. To obtain confidence intervals or test the significance of the fitted coefficients, we rely on the pivotal quantity (2.7).

3.2. *Type 3 Tobit model.* The Type 3 model is:

$$\begin{aligned}
 y_1^* &= X_1\beta_1 + \epsilon_1, \epsilon_1 \sim \mathcal{N}(0, \sigma_1^2 I) \\
 y_2^* &= X_2\beta_2 + \epsilon_2, \epsilon_1 \sim \mathcal{N}(0, \sigma_2^2 I) \\
 (3.3) \quad y_{1,i} &= \begin{cases} y_{1,i}^* & \text{if } y_{1,i}^* > 0 \\ 0 & \text{otherwise} \end{cases}, i = 1, \dots, n \\
 y_{2,i} &= \begin{cases} y_{2,i}^* & \text{if } y_{1,i}^* > 0 \\ 0 & \text{otherwise} \end{cases}, i = 1, \dots, n.
 \end{aligned}$$

The error terms ϵ_1 and ϵ_2 are again normal with covariance $\sigma_{12}^2 I$. The Type 1 (1.1) and Type 2 (3.8) Tobit models are special cases of the Type 3. Type 2 differs only in that y_1 is not observed, and Type 1 is special case of Type 2 in which $X_1 = X_2$ and $\epsilon_1 = \epsilon_2$.

To derive the two-step estimator for the Type 3 Tobit, we first evaluate the bias incurred by censoring. For y_1 , the bias has the same form as (2.1):

$$\begin{aligned}
 \mathbf{E}[y_{1,i} \mid y_{1,i}^* > 0] &= x_{1,i}^T \beta_1 + \mathbf{E}[\epsilon_{1,i} \mid \epsilon_{1,i} > -x_{1,i}^T \beta_1] \\
 &= x_{1,i}^T \beta_1 + \sigma_1 \lambda \left(\frac{1}{\sigma_1} x_{1,i}^T \beta_1 \right),
 \end{aligned}$$

where $\lambda(x)$ is the inverse Mills ratio. For y_2 , the bias is similar:

$$\mathbf{E}[y_{2,i}^* \mid y_{1,i}^* > 0] = x_{2,i}^T \beta_2 + \mathbf{E}[\epsilon_{2,i} \mid \epsilon_{1,i} > -x_{1,i}^T \beta_1].$$

Since ϵ_1 and ϵ_2 are jointly normal, $\mathbf{E}[\epsilon_{2,i} \mid \epsilon_{1,i} > -x_{1,i}^T \beta_1]$ is a simple linear function of $\mathbf{E}[\epsilon_{1,i} \mid \epsilon_{1,i} > -x_{1,i}^T \beta_1]$:

$$(3.4) \quad \mathbf{E}[y_{2,i} \mid y_{1,i}^* > 0] = x_{2,i}^T \beta_2 + \tau \lambda \left(\frac{1}{\sigma_1} x_{1,i}^T \beta_1 \right), \tau = \frac{\sigma_{12}}{\sigma_1}$$

The two-step estimator first estimates $\lambda \left(\frac{1}{\sigma_1} x_{1,i}^T \beta_1 \right)$ with probit regression and then fits a corrected linear model to uncensored observations. As long

as the probit MLE is consistent, the two-step estimator produces consistent estimates of β_1 and β_2 .¹ The corrected linear model is again heteroscedastic:

$$(3.5) \quad \text{var}[\epsilon_{2,i} \mid \epsilon_{1,i} > -x_{1,i}^T \beta_1] = \sigma_2^2 - \tau^2 (x_{1,i}^T \alpha_1 \lambda(x_{1,i}^T \alpha_1) + \lambda(x_{1,i}^T \alpha_1)^2).$$

Algorithm 2 Heckman's two-step estimator (Type 3 Tobit model)

Require: design matrix X , responses y

- 1: Estimate $\alpha_1 = \frac{1}{\sigma_1} \beta_1$ with the probit MLE:

$$\hat{\alpha}_1 = \arg \max_{\alpha} \prod_{i: y_i=0} (1 - \Phi(x_{1,i}^T \alpha)) \prod_{i: y_i \neq 0} \Phi(x_{1,i}^T \alpha).$$
 - 2: Regress the uncensored responses \bar{y}_1 on \bar{X}_1 and $\lambda(\bar{X}_1 \hat{\alpha}_1)$ to estimate β_1 :

$$\hat{\beta}_1, \hat{\sigma}_1 = \arg \min_{\beta, \sigma} \frac{1}{2} \|\bar{y}_1 - \sigma_1 \hat{\lambda} - \bar{X}_1 \beta\|_2^2.$$
 - 3: Regress the uncensored responses \bar{y}_2 on \bar{X}_2 and $\hat{\lambda} = \lambda(\bar{X}_2 \hat{\alpha}_1)$ to estimate β_2 :

$$\hat{\beta}_2, \hat{\tau} = \arg \min_{\beta, \tau} \frac{1}{2} \|\bar{y}_2 - \tau \hat{\lambda} - \bar{X}_2 \beta\|_2^2.$$
 - 4: Obtain an estimate of σ_2 from the residuals of step 3:

$$\hat{\sigma}_2 = \frac{1}{n-p} \|\bar{y}_2 - \bar{X}_2 \hat{\beta}_2\|_2^2 + \hat{\tau}^2 \frac{1}{n} \sum_{i=1}^n x_{1,i}^T \hat{\alpha}_1 \lambda(x_{1,i}^T \hat{\alpha}_1) + \lambda(x_{1,i}^T \hat{\alpha}_1)^2.$$
-

The two-step estimate of β_1 is given by the first part of

$$\hat{\gamma}_1 = \hat{Z}_1^\dagger \bar{y}_1, \quad \hat{Z}_1 = [\bar{X}_1 \quad \hat{\lambda}],$$

Given the censoring event $\bar{y}_1 > 0$, \bar{y}_1 has a constrained normal distribution. To perform (valid) post-correction inference for $\beta_{1,j} = e_j^T \bar{X}_1^\dagger \bar{X}_1 \beta_1$, we apply the results in Section 2. Let $\eta_{1,j} = (\bar{X}_1^\dagger)^T e_j$. By Corollary 2.2, we know

$$F(\eta_{1,j}^T \bar{y}_1, \beta_{1,j}, \sigma_1^2 \|\eta_{1,j}\|_2^2, \mathcal{V}_{\eta_{1,j}}^-(\bar{y}_1), \mathcal{V}_{\eta_{1,j}}^+(\bar{y}_1)) \mid \bar{y}_1 > 0 \sim \text{unif}(0, 1),$$

where $\mathcal{V}_{\eta}^+(y)$, $\mathcal{V}_{\eta}^-(y)$, $\mathcal{V}_{\eta}^0(y)$ are given by (2.3), (2.4), (2.5) with $A = -I$ and $b = 0$. To form intervals for $\beta_{1,j}$, we “invert” F to obtain

$$\mathbf{Pr}(\beta_{1,j} \in [\nu_{\alpha/2}(\bar{y}_1), \nu_{1-\alpha/2}(\bar{y}_1)] \mid \bar{y}_1 > 0) = 1 - \alpha,$$

where $\nu_{\alpha}(\bar{y})$ is the (unique) root of

$$(3.6) \quad F(\eta_{1,j}^T \bar{y}_1, \nu, \sigma_1^2 \|\eta_{1,j}\|_2^2, \mathcal{V}_{\eta_{1,j}}^-(\bar{y}_1), \mathcal{V}_{\eta_{1,j}}^+(\bar{y}_1)) = \alpha.$$

To test the significance of $\beta_{1,j}$, we form a p-value

$$F(\eta_{1,j}^T \bar{y}_1, 0, \sigma_1^2 \|\eta_{1,j}\|_2^2, \mathcal{V}_{\eta_{1,j}}^-(\bar{y}_1), \mathcal{V}_{\eta_{1,j}}^+(\bar{y}_1))$$

¹Olsen (1980) notes that the consistency of the two-step estimator does not require the joint normality of ϵ_1 and ϵ_2 provided ϵ_1 is normal and $y_{2,i}^* = x_{2,i}^T \beta_2 + \frac{\sigma_{12}}{\sigma_1^2} (y_{1,i}^* - x_{1,i}^T \beta_1) + \epsilon_{2,i}'$ for some $\epsilon_{2,i}'$ independent of $y_{1,i}^*$. The asymptotic variance of two-step estimators under these more general conditions is given by Lee (1982).

with $\text{unif}(0, 1)$ distribution under the null $H_0 : \beta_j = 0$. We reject H_0 when the p-value is smaller than $\frac{\alpha}{2}$ or larger than $1 - \frac{\alpha}{2}$. The corrected intervals for β_1 should behave comparably with the corrected intervals for β in a Type 1 Tobit model.

The two-step estimate of β_2 is given by the first part of

$$\hat{\gamma}_2 = \hat{Z}_2^\dagger \bar{y}_2, \quad \hat{Z}_2 = [\bar{X}_2 \quad \hat{\lambda}],$$

To perform post-correction inference on β_2 , we must account for the dependence between \bar{y}_1 and \bar{y}_2 . Given the censoring event $\bar{y}_1 > 0$, the pair (\bar{y}_1, \bar{y}_2) has a constrained normal distribution, i.e.

$$\begin{bmatrix} \bar{y}_1 \\ \bar{y}_2 \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} X_1 \beta_1 \\ X_2 \beta_2 \end{bmatrix}, \begin{bmatrix} \sigma_1^2 I & \sigma_{12} I \\ \sigma_{12} I & \sigma_2^2 I \end{bmatrix} \right), \text{ subject to } \bar{y}_1 > 0.$$

Let $\bar{y} = (\bar{y}_1, \bar{y}_2)$ and $\bar{\mu}, \bar{\Sigma}$ be its (unconditional) expected value and covariance. To form intervals for $\beta_{2,j}$, we first express our target as

$$\beta_{2,j} = e_j^T \bar{X}_2^\dagger \bar{X}_2 \beta_2 = \underbrace{e_j^T \begin{bmatrix} 0 & \bar{X}_2^\dagger \end{bmatrix}}_{\eta_{2,j}^T} \bar{\mu}.$$

By Corollary 2.2, we know

$$(3.7) \quad F \left(\eta_{2,j}^T \bar{y}, \beta_{2,j}, \eta_{2,j}^T \bar{\Sigma} \eta_{2,j}, \mathcal{V}_{\eta_{2,j}}^-(\bar{y}), \mathcal{V}_{\eta_{2,j}}^+(\bar{y}) \right) \mid \bar{y}_1 > 0,$$

where $\mathcal{V}_\eta^+(y)$, $\mathcal{V}_\eta^-(y)$, $\mathcal{V}_\eta^0(y)$ are given by (2.3), (2.4), (2.5), with $A = [-I \quad 0]$ and $b = 0$, is uniformly distributed. With this pivotal quantity, we derive confidence intervals and significance tests for $\beta_{2,j}$ like in Section 2.

To form confidence intervals for $\beta_{2,j}$, we “invert” F to obtain

$$\mathbf{Pr}(\beta_{2,j} \in [\nu_{\alpha/2}(\bar{y}), \nu_{1-\alpha/2}(\bar{y})] \mid \bar{y}_1 > 0) = 1 - \alpha,$$

where $\nu_\alpha(\bar{y})$ is the (unique) root of

$$F(\eta_{2,j}^T \bar{y}, \nu, \eta_{2,j}^T \bar{\Sigma} \eta_{2,j}, \mathcal{V}_{\eta_{2,j}}^-(\bar{y}), -\infty, \infty) = \alpha.$$

To test the significance of $\beta_{2,j}$, we form a p-value

$$F(\eta_{2,j}^T \bar{y}, 0, \eta_{2,j}^T \bar{\Sigma} \eta_{2,j}, -\infty, \infty)$$

with $\text{unif}(0, 1)$ distribution under the null $H_0 : \beta_{2,j} = 0$. We reject H_0 when the p-value is smaller than $\frac{\alpha}{2}$ or larger than $1 - \frac{\alpha}{2}$. Figure 5 shows result from two simulations that compare the coverage of the corrected intervals versus the normal intervals. We summarize our results in a pair of lemmas.




FIG 5. 95% corrected (left) and normal (right) intervals for $\beta_{2,1}$ in a Type 3 Tobit model over 50 simulated data sets ($n = 100$, $p_1 = p_2 = 10$, $\sigma_1^2 = \sigma_2^2 = 1.0$, $\sigma_{12} = 0.5$, standard Gaussian design). A green bar means the confidence interval covers $\beta_{2,1}$ while a red bar means otherwise. Although the corrected intervals are wider, they cover $\beta_{2,1}$ (black line) at the nominal rate. Empirically, we observe the coverage of the normal intervals becomes worse as the signal strength decreases.

LEMMA 3.1. Let $\eta_{2,j} = \begin{bmatrix} 0 & \bar{X}_2^\dagger \end{bmatrix}^T e_j$. Define $\nu_\alpha(\bar{y})$ to be the (unique) root of

$$F\left(\eta_{2,j}^T \bar{y}, \beta_{2,j}, \eta_{2,j}^T \bar{\Sigma} \eta_{2,j}, \mathcal{V}_{\eta_{2,j}}^-(\bar{y}), \mathcal{V}_{\eta_{2,j}}^+(\bar{y})\right) = \alpha.$$

Given $\bar{y} > 0$, $[\nu_{\alpha/2}(\bar{y}), \nu_{1-\alpha/2}(\bar{y})]$ is a valid $1 - \alpha$ confidence interval for $\beta_{2,j}$:

$$\Pr(\beta_{2,j} \in [\nu_{\alpha/2}(\bar{y}), \nu_{1-\alpha/2}(\bar{y})] \mid \bar{y}_1 > 0) = 1 - \alpha,$$

LEMMA 3.2. Let $\eta_{1,j} = \begin{bmatrix} 0 & \bar{X}_2^\dagger \end{bmatrix}^T e_j$. Given $\bar{y} > 0$, the test that rejects when

$$F\left(\eta_{2,j}^T \bar{y}, 0, \eta_{2,j}^T \bar{\Sigma} \eta_{2,j}, \mathcal{V}_{\eta_{2,j}}^-(\bar{y}), \mathcal{V}_{\eta_{2,j}}^+(\bar{y})\right) \in \left[0, \frac{\alpha}{2}\right) \cup \left(\frac{\alpha}{2}, 1\right]$$

is a valid α -level test for $H_0 : \beta_{2,j} = 0$.

3.3. *Type 2 Tobit model.* The Type 2 Tobit model is very similar to the Type 3 model, except the y_1^* not observed:

$$(3.8) \quad \begin{aligned} y_1^* &= X_1\beta_1 + \epsilon_1, \epsilon_1 \sim \mathcal{N}(0, I) \\ y_2^* &= X_2\beta_2 + \epsilon_2, \epsilon_2 \sim \mathcal{N}(0, \sigma_2^2 I) \\ y_{2,i} &= \begin{cases} y_{2,i}^* & \text{if } y_{1,i}^* > 0 \\ 0 & \text{otherwise} \end{cases}, i = 1, \dots, n. \end{aligned}$$

The error terms ϵ_1 and ϵ_2 are again jointly normal with covariance $\sigma_{12}^2 I$. Without loss of generality, we set $\sigma_1^2 = 1$. Type 2 is also called the *sample selection model* and plays significant parts in the evaluation of treatment effects and program evaluation. To fit a Type 2, we rely on Heckman's two-step estimator (Algorithm 2) except we skip step 2 (estimating β_1).

The two-step estimate of β_2 is given by the first part of

$$\hat{\gamma}_2 = \hat{Z}_2^\dagger \bar{y}_2, \hat{Z}_2 = [\bar{X}_2 \quad \hat{\lambda}],$$

Given the censoring event $\bar{y}_1 > 0$, the pair (\bar{y}_1, \bar{y}_2) has a constrained normal distribution, i.e. the pdf of $\bar{y} = (\bar{y}_1, \bar{y}_2)$ is

$$\begin{bmatrix} \bar{y}_1 \\ \bar{y}_2 \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} X_1\beta_1 \\ X_2\beta_2 \end{bmatrix}, \begin{bmatrix} I & \sigma_{12}I \\ \sigma_{12}I & \sigma_2^2 I \end{bmatrix} \right), \text{ subject to } \bar{y}_1 > 0.$$

Let $\bar{y} = (\bar{y}_1, \bar{y}_2)$ and $\bar{\mu}, \bar{\Sigma}$ be its (unconditional) expected value and covariance. To perform inference on $\beta_{2,j}$, we first express our target as

$$\beta_{2,j} = e_j^T \bar{X}_2^\dagger \bar{X}_2 \beta_2 = \underbrace{e_j^T \begin{bmatrix} 0 & \bar{X}_2^\dagger \end{bmatrix}}_{\eta_{2,j}^T} \bar{\mu}.$$

Since \bar{y}_1 is not observed, we cannot evaluate $\mathcal{V}_{\eta_{2,j}}^+(\bar{y}), \mathcal{V}_{\eta_{2,j}}^-(\bar{y}), \mathcal{V}_{\eta_{2,j}}^0(\bar{y})$ to form the pivotal quantity (3.7). To perform post-correction inference on $\beta_{2,j}$, we simulate the distribution of $\eta_{2,j}^T \bar{y} \mid \bar{y}_1 > 0$ with the parametric bootstrap (Efron and Tibshirani, 1994). Since the pairs $(\bar{y}_{1,i}, \bar{y}_{2,i})$ are independent, we simulate $\bar{y} \mid \bar{y}_1 > 0$ by simulating n pairs

$$\begin{bmatrix} \bar{y}_1^* \\ \bar{y}_2^* \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} x_{1,i}^T \hat{\alpha}_1 \\ x_{2,i}^T \hat{\beta}_2 \end{bmatrix}, \begin{bmatrix} 1 & \hat{\sigma}_{12} \\ \hat{\sigma}_{12} & \hat{\sigma}_2^2 \end{bmatrix} \right), \text{ subject to } \bar{y}_1^* > 0.$$

Although simulating truncated normals is, in general, intensive, the fact that \bar{y}_2^* is unconstrained, so the marginal distribution of \bar{y}_1^* is a (univariate)




FIG 6. 95% bias-corrected bootstrap (left) and normal (right) intervals for $\beta_{2,1}$ in a Type 2 Tobit model over 50 simulated data sets. The simulations are set up like the simulations shown in Figure 6. The bootstrap intervals are computed with 1000 bootstrap samples. A green bar means the confidence interval covers $\beta_{2,1}$ while a red bar means otherwise.

truncated normal, allows us to simulate $(\bar{y}_1^*, \bar{y}_2^*)$ efficiently. To simulate a pair $(\bar{y}_1^*, \bar{y}_2^*)$, we first simulate \bar{y}_1^* (with say the inverse CDF transform) and then simulate \bar{y}_2^* conditioned on \bar{y}_1^* :

$$(3.9) \quad \begin{aligned} \bar{y}_1^* &\sim TN(x_{1,i}^T \hat{\alpha}_1, 1, 0, \infty) \\ \bar{y}_2^* &\sim \mathcal{N}(x_{2,i}^T \hat{\beta}_2 + \hat{\sigma}_{12}(\bar{y}_1^* - x_{1,i}^T \hat{\alpha}_1), \hat{\sigma}_2^2 - \hat{\sigma}_{12}^2). \end{aligned}$$

Given the bootstrap distribution of $\hat{\beta}_{j,2}$, it is straightforward to obtain confidence intervals and test the significance of $\beta_{j,2}$. We refer to [Efron and Tibshirani \(1994\)](#) for details. Figure 6 shows result from two simulations that compare the coverage of bootstrap versus normal intervals. We note the bootstrap intervals are comparable in size with the normal intervals.

4. Summary and discussion. We proposed a framework for conducting post-correction inference with two-step estimators. By conditioning on the censoring event, we obtain valid confidence intervals and significance tests for the fitted coefficients. We developed the framework with the standard (Type 1) Tobit model and showed how it generalizes to handle other censored regression models.

Although we must condition on the censoring event to perform inference, the results are valid *unconditionally*. In Section 2, we really derived a family of valid intervals/tests, one per possible outcome of the censoring mechanism. Given any outcome (censoring event), as long as we form the correct

confidence interval, the interval will have the nominal coverage rate. In other words, the strategy of forming the correct interval given the censoring event inherits the validity of the conditional intervals. A similar strategy inherits the Type I error rate of the conditional tests.

Our framework for performing post-correction inference may be readily combined with the framework of Lee et al. (2013) for post-(model) selection inference. In practice, one might wish to first select a model with the data and then perform inference on the selected model. For example, one might fit a Tobit model, observe which coefficients are significant at level α , and report $1 - \alpha$ confidence intervals for the significant coefficients. However, these intervals fail to account for the randomness in the selected model and may fail to cover the target at the nominal rate. By combining our framework with the framework of Lee et al. (2013), it's possible to perform valid inference post-correction and post-(model) selection.

Acknowledgements. Will Fithian provided the geometric interpretation of Theorem 2.1. Y. Sun was partially supported by the NIH, grant U01GM102098. J.E. Taylor was supported by the NSF, grant DMS 1208857, and by the AFOSR, grant 113039.

APPENDIX

A.1. Proof of Theorem 2.1.

THEOREM A.1. *Let $y \sim \mathcal{N}(\mu, \Sigma)$. Define $a = \frac{A\Sigma\eta}{\eta^T\Sigma\eta}$ and*

$$\begin{aligned}\mathcal{V}_\eta^+(y) &= \sup_{j:a_j < 0} \frac{1}{a_j} (b_j - (Ay)_j + a_j \eta^T y) \\ \mathcal{V}_\eta^-(y) &= \inf_{j:a_j > 0} \frac{1}{a_j} (b_j - (Ay)_j + a_j \eta^T y) \\ \mathcal{V}_\eta^0(y) &= \inf_{j:a_j = 0} b_j - (Ay)_j.\end{aligned}$$

Then, $\eta^T y$ conditioned on $Ay \leq b$, $\mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y) \geq 0$, has a truncated normal distribution, i.e.

$$\eta^T y \mid Ay \leq b, \mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y) \geq 0 \sim TN(\eta^T \mu, \eta^T \Sigma \eta, \mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y)).$$

PROOF. Our proof is similar in essence to the derivation in Lee et al. (2013). First, we prove an auxiliary result that shows $Ay \leq b$ implies $\{V^-(y) \leq \eta^T y \leq V^+(y), V^0(y) \geq 0\}$ for some $\mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y)$ that are independent of $\eta^T y$.

LEMMA A.2. Let $y \sim \mathcal{N}(\mu, \Sigma)$. Define $a = \frac{A\Sigma\eta}{\eta^T\Sigma\eta}$ and

$$\begin{aligned}\mathcal{V}_\eta^+(y) &= \sup_{j:a_j < 0} \frac{1}{a_j} (b_j - (Ay)_j + a_j \eta^T y) \\ \mathcal{V}_\eta^-(y) &= \inf_{j:a_j > 0} \frac{1}{a_j} (b_j - (Ay)_j + a_j \eta^T y) \\ \mathcal{V}_\eta^0(y) &= \inf_{j:a_j = 0} b_j - (Ay)_j.\end{aligned}$$

Then, $\{Ay \leq b\}$ implies $\{\mathcal{V}_\eta^-(y) \leq \eta^T y \leq \mathcal{V}_\eta^+(y), \mathcal{V}_\eta^0(y) \geq 0\}$. Further $\mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y)$ are independent of $\eta^T y$.

PROOF. The linear constraints $Ay \leq b$ are equivalent to

$$(A.1) \quad Ay - \mathbf{E}[Ay \mid \eta^T y] \leq b - \mathbf{E}[Ay \mid \eta^T y].$$

Since conditional expectation has the form

$$\mathbf{E}[Ay \mid \eta^T y] = A\mu + a(\eta^T y - \eta^T \mu), \quad a = \frac{A\Sigma\eta}{\eta^T\Sigma\eta},$$

(A.1) simplifies to $Ay - b - a\eta^T y \leq -a\eta^T y$. Rearranging, we obtain

$$\begin{aligned}\eta^T y &\geq \frac{1}{a_j} (b_j - (Ay)_j + a_j \eta^T y) & a_j < 0 \\ \eta^T y &\leq \frac{1}{a_j} (b_j - (Ay)_j + a_j \eta^T y) & a_j > 0 \\ 0 &\leq b_j - (Ay)_j + a_j \eta^T y & a_j = 0.\end{aligned}$$

We take the sup of the lower bounds and inf of the upper bounds to deduce

$$\underbrace{\sup_{j:a_j < 0} \frac{1}{a_j} (b_j - (Ay)_j + a_j \eta^T y)}_{\mathcal{V}_\eta^-(y)} \leq \eta^T y \leq \underbrace{\inf_{j:a_j > 0} \frac{1}{a_j} (b_j - (Ay)_j + a_j \eta^T y)}_{\mathcal{V}_\eta^+(y)}.$$

Since y is normal, $b_j - (Ay)_j + a_j \eta^T y$, $j = 1, \dots, m$ are independent of $\eta^T y$. Hence $\mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y) \geq 0$ are also independent of $\eta^T y$. \square

To complete the proof of Theorem 2.1, we must show $\eta^T y$ given $Ay \leq b, \mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y) \geq 0$ is truncated normal. By Lemma A.2, the conditional law of $\eta^T y$ is

$$\begin{aligned}F(\eta^T y \mid Ay \leq b, \mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y) \geq 0) \\ = F(\eta^T y \mid \mathcal{V}_\eta^-(y) \leq \eta^T y \leq \mathcal{V}_\eta^+(y), \mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y), \mathcal{V}_\eta^0(y) \geq 0) \\ = \frac{F(\eta^T y, v^- \leq \eta^T y \leq v^+ \mid \mathcal{V}_\eta^+(y) = v^+, \mathcal{V}_\eta^-(y) = v^-, \mathcal{V}_\eta^0(y) \geq 0)}{F(v^- \leq \eta^T y \leq v^+ \mid \mathcal{V}_\eta^+(y) = v^+, \mathcal{V}_\eta^-(y) = v^-, \mathcal{V}_\eta^0(y) \geq 0)},\end{aligned}$$

Since $\mathcal{V}_\eta^+(y)$, $\mathcal{V}_\eta^-(y)$, $\mathcal{V}_\eta^0(y)$ are independent of $\eta^T y$,

$$\begin{aligned} F(\eta^T y \mid Ay \leq b, \mathcal{V}_\eta^+(y) = v^+, \mathcal{V}_\eta^-(y) = v^-, \mathcal{V}_\eta^0(y) \geq 0) \\ = \frac{F(\eta^T y, v^- \leq \eta^T y \leq v^+)}{F(v^- \leq \eta \leq v^+)}, \end{aligned}$$

which is the CDF of $TN(\eta^T \mu, \eta^T \Sigma \eta, \mathcal{V}_\eta^+(y), \mathcal{V}_\eta^-(y))$ □

A.2. Monotonicity of F .

LEMMA A.3. *Let $F(x, \mu, \sigma^2, a, b)$ be the CDF of a truncated normal random variable. Then F is monotone decreasing in μ .*

PROOF. The truncated normal distribution is a natural exponential family in the mean μ . Thus, its likelihood ratio is monotone in μ , i.e. for $\mu_1 < \mu_2$ and $y_2 < y_1$,

$$\frac{f_{\mu_2}(y_1)}{f_{\mu_2}(y_2)} < \frac{f_{\mu_1}(y_1)}{f_{\mu_1}(y_2)}.$$

This implies $f_{\mu_2}(y_1)f_{\mu_1}(y_2) < f_{\mu_2}(y_2)f_{\mu_1}(y_1)$. We integrate with respect to y_1 over $(-\infty, y]$ and with respect to y_2 over $[y, \infty)$ to obtain

$$(1 - F_{\mu_1}(y))F_{\mu_2}(y) < (1 - F_{\mu_2}(y))F_{\mu_1}(y).$$

We subtract the cross-terms to conclude $F_{\mu_2}(y) < F_{\mu_1}(y)$. □

REFERENCES

- AMEMIYA, T. (1973). Regression analysis when the dependent variable is truncated normal. *Econometrica: Journal of the Econometric Society* 997–1016.
- AMEMIYA, T. (1985). *Advanced Econometrics*. Harvard University Press.
- DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 1–38.
- EFRON, B. and TIBSHIRANI, R. J. (1994). *An Introduction to the Bootstrap* **57**. CRC press.
- GOLDBERGER, A. S. (1964). *Econometric Theory*. New York: John Wiley & Sons.
- HECKMAN, J. J. (1976). The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models. *Annals of Economic and Social Measurement, Volume 5, number 4* 475–492.
- HECKMAN, J. J. (1979). Sample selection bias as a specification error. *Econometrica: Journal of the econometric society* 153–161.
- KEELEY, M. C., ROBINS, P. K., SPIEGELMAN, R. G. and WEST, R. W. (1978). The Estimation of Labor Supply Models Using Experimental Data. *The American Economic Review* **68** pp. 873–887.
- LEE, L.-F. (1982). Some approaches to the correction of selectivity bias. *The Review of Economic Studies* **49** 355–372.

- LEE, J. D., SUN, D. L., SUN, Y. and TAYLOR, J. E. (2013). Exact post-selection inference with the lasso. *arXiv preprint arXiv:1311.6238*.
- OLSEN, R. J. (1978). Note on the uniqueness of the maximum likelihood estimator for the Tobit model. *Econometrica: Journal of the Econometric Society* 1211–1215.
- OLSEN, R. J. (1980). A least squares correction for selectivity bias. *Econometrica: Journal of the Econometric Society* 1815–1820.
- POWELL, J. L. (1984). Least absolute deviations estimation for the censored regression model. *Journal of Econometrics* **25** 303–325.
- PUHANI, P. (2000). The Heckman correction for sample selection and its critique. *Journal of Economic Surveys* **14** 53–68.
- TOBIN, J. (1958). Estimation of relationships for limited dependent variables. *Econometrica: Journal of the Econometric Society* 24–36.
- WITTE, A. D. (1980). Estimating the economic model of crime with individual data. *The Quarterly Journal of Economics* **94** 57–84.

INSTITUTE FOR COMPUTATIONAL AND MATHEMATICAL ENGINEERING
STANFORD UNIVERSITY
475 VIA ORTEGA, STANFORD, CA 94305
E-MAIL: yuekai@stanford.edu

DEPARTMENT OF STATISTICS
STANFORD UNIVERSITY
390 SERRA MALL, STANFORD, CA 94305
E-MAIL: jonathan.taylor@stanford.edu